Seasonal Evolution of Influenza-Related Mortality

Todd Graves* **and Rick Picard**

Los Alamos National Laboratory

Los Alamos, New Mexico 87545

**email:* tgraves@lanl.gov

SUMMARY. We hierarchically model pneumonia and influenza mortality as a function of time, capturing the onset, duration, and severity of a flu season. Using historical data, we fit a model in which weekly mortality has an overdispersed Poisson distribution. This approach allows us to deal directly with year-to-year variation and to characterize that variation over longer time periods. Results from this analysis challenge the conventional wisdom regarding the extent to which influenza epidemics have occurred in the U.S. during the past 40 years.

KEY WORDS: biosurveillance; forecasting; hierarchical modeling; time series.

1. Introduction

Pneumonia and influenza (P&I) mortality is one measure of the severity of a flu season. P&I data since 1996 are available in National Notifiable Diseases Surveillance System (NNDSS) "Provisional Weekly Tables" at the www.cdc.gov/scientific.htm web site. For many cities, mortality has been monitored since 1962 and has been frequently examined for many purposes, such as estimating "excess deaths" attributable to influenza.

Our purpose is to model the seasonal nature of P&I mortality in order to understand the past and to predict the (near term) future. In particular, given information from an ongoing flu season, we predict deaths for the remainder of the season and assign uncertainties to those predictions. An accurate forecasting methodology is useful for many purposes, including response planning for anticipated near term levels of influenza and biosurveillance to detect potential anomalies.

At the core of good prediction is an explicit accounting for season-specific behavior. To illustrate, consider Figure 1, which displays weekly P&I mortality from three flu seasons in Los Angeles (the overlaid curves in Figure 1 are discussed later). Year-to-year variability is apparent: the peak mortality levels and their corresponding times of occurrence range from more than 60 deaths near week 30 for the 1968-9 season to roughly 40 deaths near week 38 for 1975-6 to less than 20 deaths with no discernible peak time for 1973-4 (we define week $t = 0$ to be the 27^{th} week of the calendar year to center the plots).

The season-to-season variability reflects that some flu strains are more virulent than

others and that, in a given city, flu may arrive at times that are comparatively early or late in the season. As a flu season evolves, we estimate season-specific parameters that reflect a baseline level, peak height, and the time at which the peak occurs. Related parameter estimates, derived from a model assuming a slowly moving baseline coupled with a Gaussian-shaped peak and an overdispersed Poisson error term, are then used in prediction.

Beyond the issue of prediction, parameter estimates covering the past 40 years provide insight into long term behavior of influenza-related mortality. That insight leads us to advance the potentially controversial assertion that the extent to which epidemics have occurred is greatly exaggerated.

In Section 2, we briefly review previous modeling of P&I mortality data and present the hierarchical approach. We then apply the hierarchical model to P&I data in Section 3 and illustrate its predictive capability. A variety of related remarks are given in Section 4.

2. Methods

2.1 Cyclical Regression and Time Series Models

The most well known modeling of P&I data is displayed in Figure 2. Simonsen et. al. (1997a) nicely summarize the methodology for cyclical regression. Related applications (Lui and Kendal 1985; Simonsen et. al. 1997b; Simonsen et. al. 2000; Mostashari 2002) have also appeared.

The cyclical regression model has the form

$$Y(t) \;=\; a + bt + ct^2 + d\,cos\left(\frac{2\pi t}{52.167}\right) + f\,sin\left(\frac{2\pi t}{52.167}\right) + e_t\;, \tag{1}$$

where $Y(t)$ is the percentage of deaths attributed to P&I in week t, 52.167 is the average number of weeks in a year, (a, b, c, d, f) are model parameters, and e_t is an error term.

Several aspects of cyclical regression warrant attention, as the related issues must be dealt with in any practical modeling. The first involves normalization: during the 40+ year monitoring period, the U.S. population has increased by more than 50%. As such, P&I deaths, when measured in absolute number, increase over time for many metropolitan areas owing to the size and age of the population at risk. Cyclical regression accounts for this by normalizing P&I deaths to the total number of deaths.

A second aspect of cyclical regression involves its characterization of baseline behavior. In the absence of seasonal effects (i.e., set the parameters d and f in (1) equal to zero), the model postulates that P&I deaths behave according to a slowly moving baseline, described by a quadratic function $a + bt + ct^2$ in time t over the relevant time period.

Thirdly, note in Figure 2 that P&I results are accumulated over all nationwide surveillance sites. By combining information across many cities, data from those cities where flu arrives early in a season are effectively offset by results from other cities where it arrives later. We return to this point in the next section.

The final aspect of cyclical regression involves the error bands in Figure 2. These error bands are obtained by fitting the data "after excluding weeks with excess mortality" (Simonsen et. al. 1997a, p. 391) and then computing baseline uncertainties from the remaining data only. At that point, other data are superimposed on the fit and excesses are revealed.

Conventional time series modeling can also be used on P&I data. Good references include Stroup, Thacker, and Herndon (1988), and Williamson and Hudson (1999), so we do not detail that work here. Upon coupling ARIMA modeling with statistical process control techniques for one-step-ahead forecast errors, aberrations are detected in near real time. The techniques used are similar to those for similar problems in industrial manufacturing (Montgomery and Mastrangelo 1991) and in nuclear materials safeguards (Picard 1987). Such detection methods could also be useful for biosurveillance.

2.2 Hierarchical Modeling of Historical Data

As is apparent from Figure 1, applying models such as those in the previous section to city-specific data is ill-advised. Procedures that force-fit a one-size-fits-all peak height, peak width, and time of year at which the peak occurs, are poorly suited to cases where season-to-season variability is important. Our goal is to overcome this model deficiency and develop a predictive mechanism for P&I mortality.

The model we use combines a slowly moving baseline term together with a Gaussian-shaped peak to capture seasonal effects. To begin, let $E[M_s(t)]$ denote the expected P&I

mortality for week t of season s, where

$$E[M_s(t)] \;=\; b_s(t) \;+\; \frac{c_s}{\sigma_s}\,\phi\left(\frac{t-\Delta_s}{\sigma_s}\right)\;, \tag{2}$$

$b_s(t)$ is the baseline level, c_s quantifies the contribution of season s to overall mortality (as discussed below), and the function $\phi(\cdot)$ is the standard normal density. The Gaussian function $\phi(\cdot)$ is indexed by easily interpreted parameters Δ_s, denoting the time at which the peak occurs, and σ_s, which corresponds to the duration of flu season s. Curves of the form (2) are overlaid on Figure 1.

The baseline function has the form

$$b_s(t) \;=\; b_{s-1} \;+\; \frac{t}{52}(b_s - b_{s-1}) \tag{3}$$

for $\{b_s\}$ the mortality levels for week $t = 0$ of the individual flu seasons. In other words, the slowly moving baseline is modeled by linearly interpolating between annual off-peak baseline levels. Were the baseline values $\{b_s\}$ to conform to a quadratic function and were there no season-to-season differences (set c_s, Δ_s and σ_s to the same values for all flu seasons), then the hierarchical model would be nearly indistinguishable from cyclical regression.

The expected number of P&I deaths for flu season s is obtained by integrating the

expected mortality $E[M_s(t)]$ over time, getting

$$\int_{t=0}^{52} E[M_s(t)]\,dt \;=\; c_s \;+\; 52 \;\times\; \frac{b_{s-1}+b_s}{2}\;, \tag{4}$$

which is simply the sum of the seasonal component and a baseline contribution.

To continue the model specification, a stochastic error term is needed. Exploratory analyses indicated that the data were overdispersed relative to the Poisson distribution. Overdispersion here may arise from several causes: for example, when there is variation in host susceptibility, the population propensity to die from P&I over time/space varies around a population average. Also, the reporting system evolves over time: during the roughly 40 years of data collection, changes have occurred in reporting practices, in the people involved with the reporting, in various procedures (e.g., the way that causes of death are categorized and the proportion of deaths entailing autopsies), and so on.

A common model for such overdispersion is the negative binomial (McCullagh and Nelder 1989, p. 199). The corresponding density has the form

$$\text{Prob}\,\{M_s(t)=x_s(t)\,|\,\mu=E[M_s(t)]\} \;=\; \frac{\Gamma(x_s(t)\;+\;\psi\,\mu)\;\;\psi^{\psi\,\mu}}{x_s(t)!\;\Gamma(\psi\,\mu)\;(1+\psi)^{x_s(t)+\psi\;\mu}}\;. \tag{5}$$

Here, the parameter $\psi > 0$ inflates the variance:

$$\text{Var}\{M_s(t)\,|\,\mu=E[M_s(t)]\} \;=\; \frac{1+\psi}{\psi}\,\mu\,, \tag{6}$$

and the limiting case ($\psi \to \infty$) is a Poisson density, having variance μ.

To complete the hierarchical specification, we capture the historical behavior of the season-to-season parameters using prior distributions. Over the 40-year time period of the mortality data, the baseline values $\{b_s\}$, seasonal contributions $\{c_s\}$, and times $\{\Delta_s\}$ at which the peaks occur often behave in a statistically predictable fashion. As we illustrate using data from Albuquerque, the simplest form of the model postulates parameters that behave from season to season as

$$\{\log b_s\} \sim i.i.d.\ N(\mu_b, \sigma_b^2)\ , \tag{7}$$

$$\{\log c_s\} \sim i.i.d.\ N(\mu_c, \sigma_c^2)\ , \tag{8}$$

$$\{\Delta_s\} \sim i.i.d.\ N(\mu_\Delta, \sigma_\Delta^2)\ , \tag{9}$$

$$\{\log \sigma_s\} \sim i.i.d.\ N(\mu_\sigma, \sigma_\sigma^2)\ , \text{ and} \tag{10}$$

$$\psi \sim \mathrm{Ga}(\mu_\psi, \sigma_\psi^2)\ , \tag{11}$$

where the gamma distribution for ψ is parameterized by its mean and variance. The model does not assume a correlation between the peak mortality level and the time in the flu season at which it occurs. Also, it postulates that, except for an inherited baseline term and a common overdispersion parameter, there is no time dependency from one flu season to the next. Finally, the hyperparameters μ_c, σ_c, and so forth are also assigned prior distributions, completing the hierarchical specification.

Obvious generalizations of the above steady-state hierarchical modeling assumptions can be pursued when departures from that situation are warranted. A simple example of this involves metropolitan areas for which the population at risk changes significantly during the time frame covered by the modeling. Modifications of the *i.i.d.* nature of the baseline levels $\{b_s\}$, such as to allow those levels to change commensurate with changes in the population at risk, could be explicitly accounted for.

The normal density is used in (2) to capture the shape of the annual mortality profile. Parameters defining the normal have natural interpretations in this context, identifying the time of peak mortality and the duration of the flu season. Residual analysis for this choice of shape function is discussed later.

2.3 Estimation

We detail the calculations for obtaining season-specific parameter estimates and for obtaining predicted values. Let $\mathbf{x}$ be shorthand notation for the set of all observed mortality values $\{x_s(t)\}$, where that set includes all 52 weeks t for all seasons s. And let $\mathbf{\Theta}$ denote the set of all model parameters for all flu seasons.

Upon multiplying the prior density for $\mathbf{\Theta}$ implicit in (7)-(11) by the sampling distribution for $\{x_s(t)\}$ implicit in (5), the posterior density for $\mathbf{\Theta}$ is proportional to

$$f(\,\mathbf{\Theta}|\mathbf{x}\,) \quad \propto \quad p(\,\mathbf{\Theta}\,) \;\times\; \text{Prob}\,[M_s(t) = x_s(t)\,;\; s = 1, \ldots, S,\, t = 1, \ldots, 52\,|\,\mathbf{\Theta}\,]$$

$$
\begin{aligned}
= \quad & p_h(\mathbf{\Theta}_h) \prod_{s=1}^{S} \{\, p(\,\mathbf{\Theta}_s\,) \;\times\; \mathrm{Prob}\,[M_s(t) = x_s(t)\,;\, t = 1, \ldots, 52 \,|\, E[M_s(t)], \psi]\} \\
= \quad & p_h(\mathbf{\Theta}_h) \;\times\; \frac{1}{\sigma_b\sqrt{2\pi}}\, e^{(log(b_0)-\mu_b)^2/2\sigma_b^2} \;\times\; \frac{(\mu_\psi\sigma_\psi^2)^{\mu_\psi^2\sigma_\psi^2}\; \psi^{\mu_\psi^2\sigma_\psi^2-1}\, e^{-\mu_\psi\sigma_\psi^2\psi}}{\Gamma(\,\mu_\psi^2\,\sigma_\psi^2\,)} \\
& \times \prod_{s=1}^{S} \left\{ \frac{1}{\sigma_c\sqrt{2\pi}}\, e^{(log(c_s)-\mu_c)^2/2\sigma_c^2} \times \frac{1}{\sigma_b\sqrt{2\pi}}\, e^{(log(b_s)-\mu_b)^2/2\sigma_b^2} \right. \\
& \times \frac{1}{\sigma_\Delta\sqrt{2\pi}}\, e^{(\Delta_s-\mu_\Delta)^2/2\sigma_\Delta^2} \times \frac{1}{\sigma_\sigma\sqrt{2\pi}}\, e^{(log(\sigma_s)-\mu_\sigma)^2/2\sigma_\sigma^2} \\
& \times \left. \prod_{t=1}^{52} \frac{\Gamma(x_s(t) \,+\, \psi\, E[M_s(t)])\;\; \psi^{\psi\, E[M_s(t)]}}{x_s(t)!\;\Gamma(\psi\, E[M_s(t)])\;(1+\psi)^{x_s(t)+\psi\; E[M_s(t)]}} \right\}
\end{aligned}
$$

Here $\mathbf{\Theta}_h$ is the vector of hyperparameters defining the distributions in (7)-(11) and $p_h(\mathbf{\Theta}_h)$ is the hyperprior, formed by combining normal hyperpriors for the means and gamma hyperpriors for the standard deviations. Means and variances of those hyperpriors were obtained from the posterior distribution of city-specific parameters in another hierarchical analysis of data from several cities. For completeness, we note that $E[M_s(t)]$ is also a function of unknown parameters $\mathbf{\Theta}_s$ as per (2), so that $f(\,\mathbf{\Theta}\,|\,\mathbf{x}\,)$ is algebraically messy.

It is impractical to pursue analytical results for messy functions of more than 150 parameters, as for Albuquerque's historical flu seasons. Instead, results are obtained via Markov chain Monte Carlo, or "MCMC" (e.g., Besag, et. al. 1995). We used a general-purpose system (Graves 2001) which emphasizes Metropolis and Metropolis-Hastings moves to individual parameters, and also makes it easy to improve MCMC convergence properties by proposing Metropolis moves to multiple correlated parameters simultaneously. Many such

moves were required in this analysis, owing to high correlations between related parameters at different levels in the hierarchical model.

2.4 Prediction Methodology

We focus on two prediction problems:

1) given P&I data from an ongoing flu season, predict the next week's mortality and assign an uncertainty to that prediction, and

2) given P&I data from an ongoing flu season, predict the total mortality for the entire season and again assign an uncertainty.

To construct predictions for future weeks using data from an ongoing flu season, let $\mathbf{x}_s(t)$ denote the mortality values observed up to and including week t of season s. Suppose that mortality data $\mathbf{x}_s(t)$ have been observed and it is of interest to predict the mortality $M_s(t+1)$ for week $t+1$ of season s. There are several ways to carry out the MCMC, the simplest of which involves simulating a sample from the predictive distribution.

To that end, each member from the MCMC sample for the posterior $g[\mathbf{\Theta}_s \,|\, \mathbf{x}_s(t)]$ is used to compute $E[M_s(t+1) \,|\, \mathbf{\Theta}_s]$. For the nth such value, $\mathbf{\Theta}_s^{(n)}$, we simulate $\gamma^{(n)}$ from a gamma distribution having mean $E[M_s(t+1) \,|\, \mathbf{\Theta}_s^{(n)}]$ and index $\psi\, E[M_s(t+1) \,|\, \mathbf{\Theta}_s^{(n)}]$, and then generate $M_s(t+1)^{(n)}$ as a Poisson random variable with mean $\gamma^{(n)}$. The set $\{M_s(t+1)^{(n)}\}$ so obtained is an MCMC sample from the predictive distribution. The predictive mean

is estimated by the average of the $\{E[M_s(t+1) \,|\, \boldsymbol{\Theta}_s^{(n)}]\}$. Posterior predictive intervals are obtained from the quantiles of the simulated predictive distribution.

Should the current week of the flu season be week $t = 51$, then the above predictive distribution summarizes the remainder of the season's mortality. Otherwise, prediction of the residual mortality for the remaining $52 - t$ weeks of the season amounts to obtaining the predictive distribution of the $\{M_s(t+j),\ j = 1, \ldots, (52-t)\}$ yet to be observed. This is done by taking each member of the MCMC sample for the posterior and generating a set of independent $\{\gamma_j\}$ from gamma distributions with respective means $\{E[M_s(t+j) \,|\, \boldsymbol{\Theta}_s]\}$ and respective indices $\{\psi\, E[M_s(t+j) \,|\, \boldsymbol{\Theta}_s\}$. Then the $\{M_s(t+j),\ j = 1, \ldots, 52-t\}$ are simulated as independent Poisson random variables with means $\{\gamma_j\}$. The residual mortality, equal to the sum of the $\{M_s(t+j)\}$, is then computed. Accumulating such sums for each member of the MCMC sample yields the predictive distribution of the residual mortality.

3. Results

3.1 Application to Historical Albuquerque Data

We illustrate application of the model to metropolitan Albuquerque, a medium-sized area of interest to us because of its local proximity. Estimates and standard deviations of the hyperparameters in (7)-(11) specific to Albuquerque are

$$(\mu_b, \sigma_b) \;\; = \;\; (1.1, 0.3)\ , \qquad (\mu_c, \sigma_c) \;=\; (3.8, 0.4)\ ,$$

$$
\begin{aligned}
(\mu_\Delta, \sigma_\Delta) &= (32.3, 3.8)\ , \quad (\mu_\sigma, \sigma_\sigma) = (1.8, 0.5)\ , \text{ and} \\
(\mu_\psi, \sigma_\psi) &= (2.3, 0.2)\ . \qquad (12)
\end{aligned}
$$

The hyperparameter $\mu_\Delta \approx 32.3$ indicates that, historically, the peak mortality occurs on average in early February. And the overdispersion as per (6) inflates variability by a factor of roughly $(1 + \psi)/\psi \approx 1.4$ relative to that from a simple Poisson model. We use these hyperparameter estimates for prediction in Section 3.2.

Time sequence plots of estimated season-to-season parameters are given Figure 3 and indicate no discernible long term time trend in mortality. A common speculation is that the seasonal contributions $\{c_s\}$ should evolve slowly over time, reflecting natural evolution in the prevailing influenza strains. A diagnostic check of autocorrelations, however, provides no reason to reject the model's assumed season-to-season independence of those terms. Figure 4 displays normal quantile-quantile plots of the estimates and shows that the postulated normal distributions in (7)-(9) reasonably describe empirical behavior. Also, the plot of the $\{\log c_s\}$ against the $\{\Delta_s\}$ indicates no apparent relation between the overall severity of a flu season and the time of year at which the peak mortality occurs (Figure 5).

Residual analysis reveals minor imperfections in the model. As might be gleaned from Figure 1, pre-peak mortality rises slightly more quickly than the post-peak mortality declines, meaning that symmetric shapes such as the normal density in (2) and the sine and cosine harmonics in (1) do not fully capture the seasonal P&I cycle. But despite these minor

imperfections, we conclude from the usual model diagnostics that the steady-state model of the previous section, although purely empirical, provides a good fit to the 1962-1999 Albuquerque P&I mortality data and is useful for prediction.

Similar goodness of fit checks reveal that the steady-state model requires modification for certain other cities and time frames (see Section 4.2).

3.2 Prediction for Albuquerque

To illustrate prediction, we mimic the use of the model in a real-time environment using data from Albuquerque. Using the 36 flu seasons from mid-1962 through mid-1998 to define the initial observed data $\mathbf{x}_s(t)$, we predict the mortality $M_{s+1}(1)$ for the first week of flu season $s+1$ (1998-9) as well as the total mortality for that season.

Next, the clock is advanced one week, so that the observed data now include results from the first week of season $s+1$. Mortality for the second week is predicted, as well as the residual mortality for the season. The clock is then advanced another week, and so on.

Figure 6 summarizes the results. The one-step-ahead forecasts are plotted in the upper graph along with their corresponding 95% uncertainty intervals. The lower graph is analogous to the upper graph, dealing with the season's residual mortality instead of the one-step-ahead forecasts. In both cases, widths of the uncertainty intervals are commensurate with the size of the quantity being predicted.

Overlaid on Figure 6 are the actual results. Note that the coverage probabilities of

one-step ahead forecasts for discrete distributions tend to slightly exceed the nominal ones (the interval endpoints being integer valued means that the corresponding probabilities will slightly exceed 95%) and that there are correlations among forecast errors of residual mortalities in the same season. As such, the coverages in Figure 6 are in accordance with expectations, which is another goodness of fit check of the model.

The hierarchical approach could be applied to other metrics sensitive to the severity of a flu season, such as age-specific P&I mortality (Stroup et. al. 1988), the number of positive influenza isolates in laboratory samples (Mostashari 2002), and the proportion of patient visits to sentinel physicians for influenza-like illness, which state health departments such as New Mexico's monitor as one means of tracking an ongoing flu season. Such open-ended efforts remain as future research.

4. Discussion

4.1 Do Flu Epidemics Truly Exist?

Some flu seasons are unquestionably worse than others. But at this point we digress to consider whether epidemics have occurred during the past 40 years.

To that end, the immediate semantic question is: exactly what constitutes an "epidemic?" For the moment, consider a working definition analogous to that for an outlier. In other words, if one or more data points appear to arise from a different statistical population from that for the remainder of the data, then the data point(s) in question are deemed outliers.

It follows from this definition, for example, that simply being the most extreme value among a smooth continuum of values is not sufficient to be regarded as an outlier.

Using this definition, there appear to be no epidemic flu seasons for the P&I data we examined. In Figure 4, the normal quantile-quantile plot of estimated log seasonal contributions to mortality for Albuquerque does not appear to contain outliers. Quantile-quantile plots for the $\{\log c_s\}$ of the six other cities we examined are similarly well behaved.

As such, we conclude that data from flu seasons with comparatively high P&I mortality do *not* reflect anomalous activity, but instead conform to the upper tail of a predictable statistical population. For the record, we note that the P&I data go back only to 1962, thus excluding flu seasons such as the famous 1918 pandemic. Nor did we examine results from all 122 cities in the surveillance network, just seven. Nor did we assess influenza-related metrics other than mortality. Nonetheless, this single-population description for the ensemble of seasonal mortalities contrasts sharply with the two-population (epidemic and non-epidemic) description of flu seasons often given in the literature, where epidemics are portrayed as commonly occurring, such as in Figure 2.

4.2 Data Quality

Analyses are necessarily limited by their reliance on reported data. As has been noted, there are "special features of public heath surveillance data . . . including that the data are not usually generated from a random sample," which present "an analytic challenge" (Williamson

and Hudson 1999, p. 3284). One such challenge involves ensuring that mortality data are truly comparable over time.

Our example data set (Albuquerque) illustrates the issue. The 1999-2000, 2000-2001, and 2001-2002 flu seasons involved more than 400 P&I deaths each, compared with a seasonal maximum of 337 deaths during 1962-1999. Considering that those flu seasons did not produce unusual P&I mortality nationally, the sudden Albuquerque increase is suspicious. In actuality, an influx of biosurveillance funding together with an October 1999 revision to the way death certificate information was used to categorize P&I deaths led to changes in the monitoring system that complicate analysis of the most recent data.

More generally, hierarchical modeling must be done with care. For example, baseline parameter estimates for Des Moines abruptly change around 1982, while those for Charlotte consistently increase over time; see the goodness of fit checks in Figure 7. Obvious modifications to the model should be considered when steady-state assumptions are unrealistic.

4.3 Discussion

Hierarchical modeling is useful for many purposes, the most obvious ones being discussed in Section 3. Historical P&I behavior can sometimes be nicely summarized, such as by the steady-state parameters in (12) for Albuquerque.

But there are other uses of hierarchical modeling, such as the prompt identification of anomalous events. In Figure 6, for example, the uncertainty bands quantify the anticipated

range of mortality under normal conditions. As such, those uncertainty bands provide a signal-to-noise context for understanding what magnitudes of anomalous events would be detectable under what conditions (note that the noise level varies depending on the particular flu season involved and on the time of year). Granted, other health metrics are better suited for early warning purposes than are mortality data, but real-time monitoring procedures based on the equivalent of Figure 6 are useful for biosurveillance.

Still another use of hierarchical modeling involves the identification of phenomena warranting further investigation. For example, peak P&I mortality in New York occurs three and a half weeks earlier than in Los Angeles, on the average (see Figure 8). In turn, Los Angeles flu seasons are consistently shorter than those in Pittsburgh, as quantified by estimates of the parameter σ in the hierarchical model (2). We resist speculation regarding cause-and-effect explanations for such phenomena, but suspect that additional research could lead to a better understanding of P&I mortality than currently exists.

REFERENCES

Berger, J. O. (1989). *Statistical Decision Theory and Bayesian Analysis*. Springer-Verlag: New York.

Besag, J., Green, P., Higdon, D. and Mengersen, K. (1995). Bayesian computation and stochastic systems. *Statistical Science* **10**, 3-41.

Graves, T. L. (2001). YADAS: An object-oriented framework for data analysis using

Markov chain Monte Carlo. Los Alamos National Laboratory Technical Report LA-UR-01-4804.

Lui, K. J. and Kendal, A. P. (1987). Impact of influenza epidemics on mortality in the United States from October 1972 to May 1985. *American Journal of Public Health* **77**, 712-716.

McCullagh, P. and Nelder, J. A. (1989). *Generalized Linear Models*. Chapman Hall: London.

Montgomery, D. C. and Mastrangelo, C. M. (1991). Some statistical process control methods for autocorrelated data. *Journal of Quality Technology* **23**, 179-204.

Mostashari, F. (2002). Syndromic Surveillance in New York City. Proceedings: National Syndromic Surveillance Conference.

Picard, R. R. (1987). Sequential analysis of materials balances. *Journal of the Institute of Nuclear Materials Management* **15**, 38-42.

Simonsen L., Clarke, M. J., Stroup, D. F., Williamson, G.D., Arden, N. H., and Cox, N. J. (1997a). A method for timely assessment of influenza-related mortality in the United States. *Epidemiology* **8**, 390-395.

Simonsen, L., Clarke, M. J., Williamson, G. D., Stroup, D. F., Arden, N. H., and Schonberger, L. B. (1997b). The impact of influenza epidemics on mortality: introducing a severity index. *American Journal of Public Health* **87**, 1944-1950.

Simonsen, L., Fakuda, K., Schonberger, L. B., and Cox, N. J. (2000). The impact of influenza epidemics on hopitalizations. *Journal of Infectious Diseases* **181**, 831-837.

Stroup, D. F., Thacker, S. B., and Herndon, J. L. (1988). Application of multiple time-series analysis to the estimation of pneumonia and influenza mortality by age (1962-1983). *Statistics in Medicine* **7**, 1045-1059.

Williamson, G. D. and Hudson G. W. (1999). A monitoring system for detecting aberrations in public health surveillance reports. *Statistics in Medicine* **18**, 3283-3298.

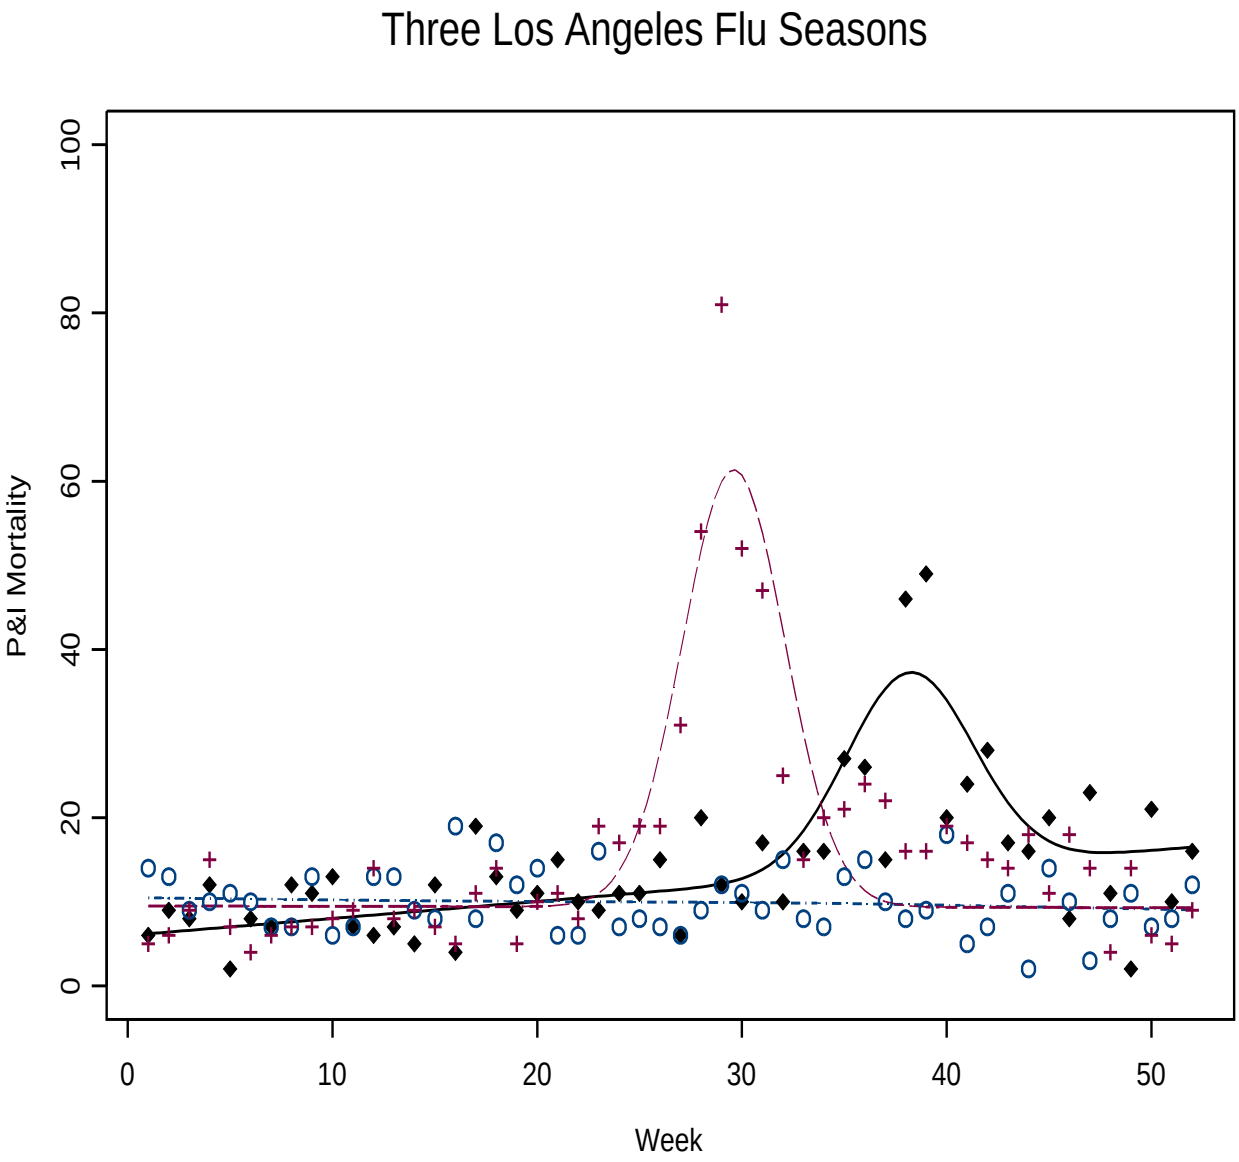

Figure 1: Weekly P&I Mortality in Los Angeles for three flu seasons: 1968-9 (denoted by a '+' symbol), 1973-4 (denoted by an open circle) and 1975-6 (denoted by a solid diamond).

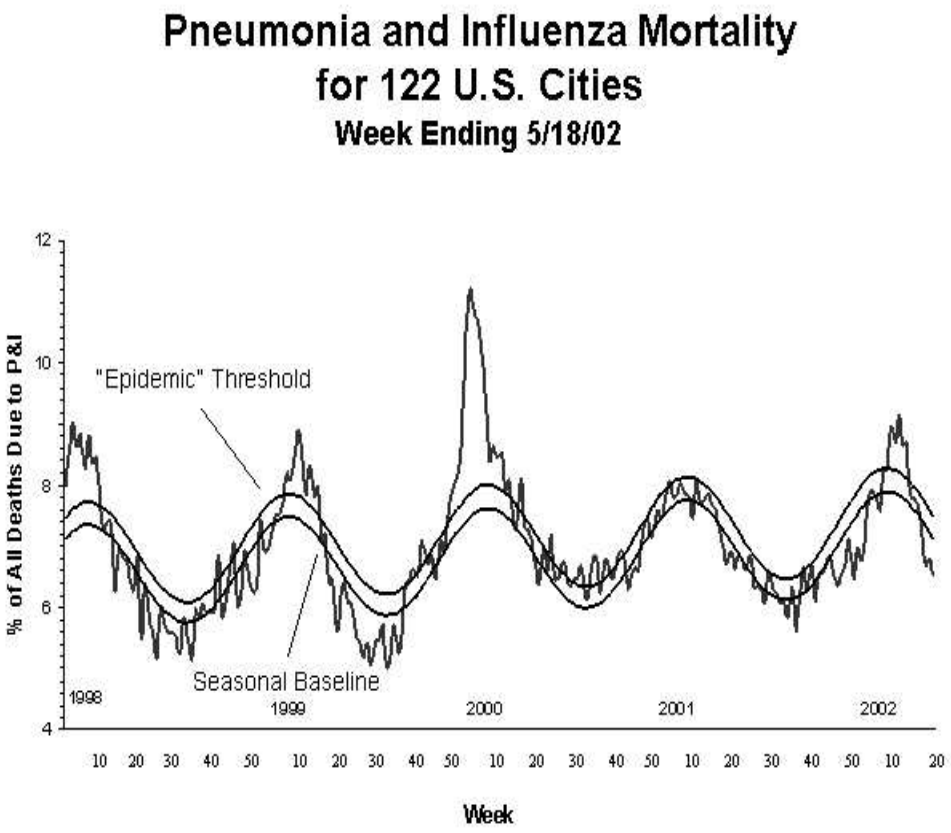

Figure 2: Cyclical Regression Model for P&I Data, Reproduced From the CDC Web Site.

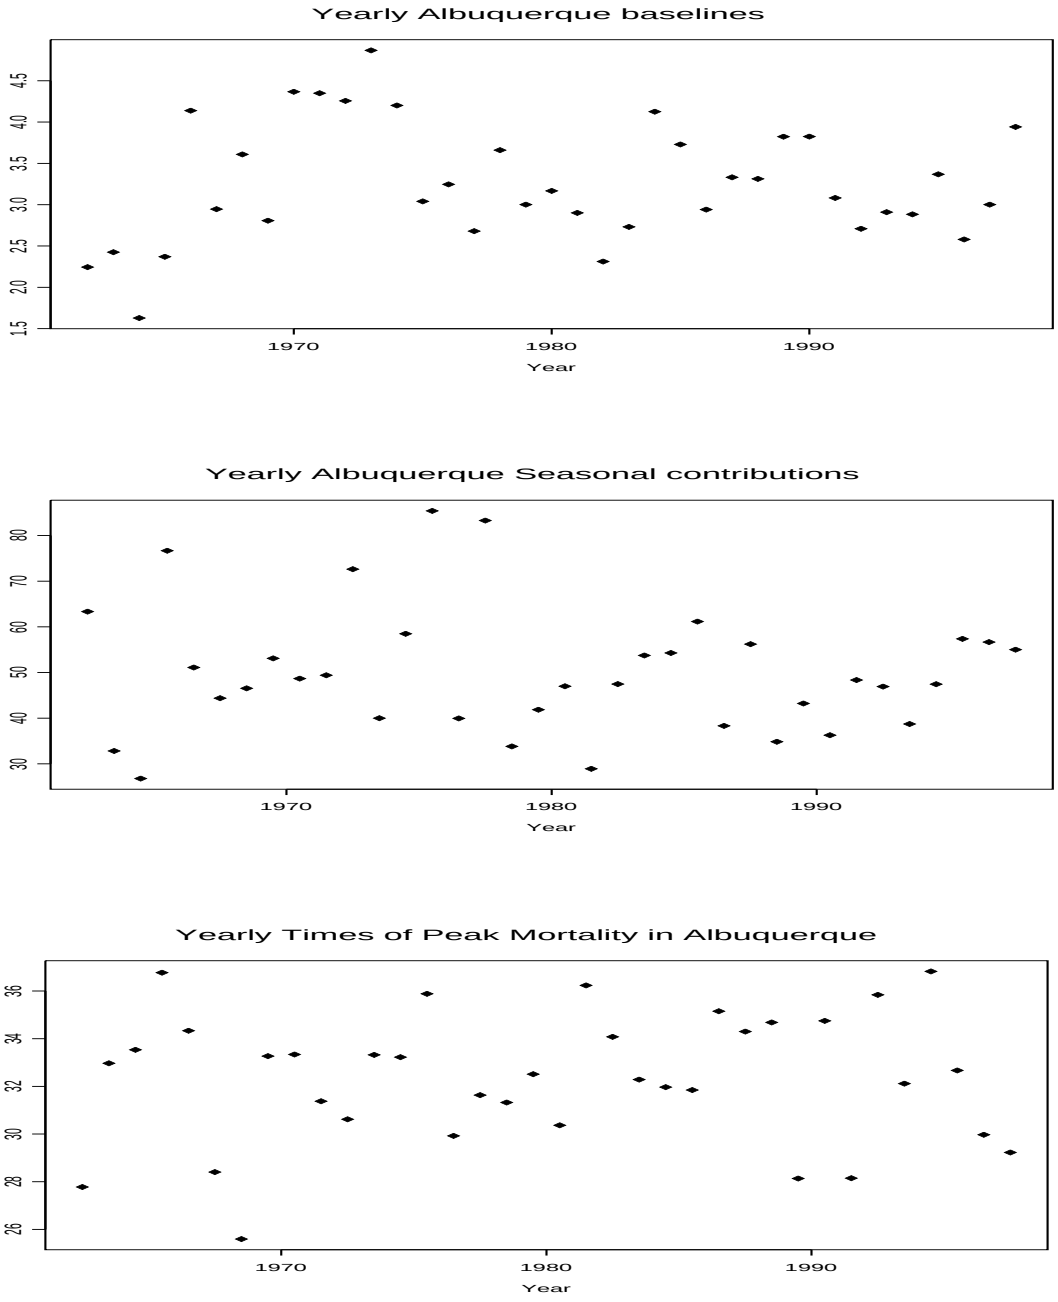

Figure 3: Time Sequence Plots of Estimated Baselines $\{b_s\}$, Seasonal Contributions $\{c_s\}$, and Times $\{\Delta_s\}$ of Peak Activity.

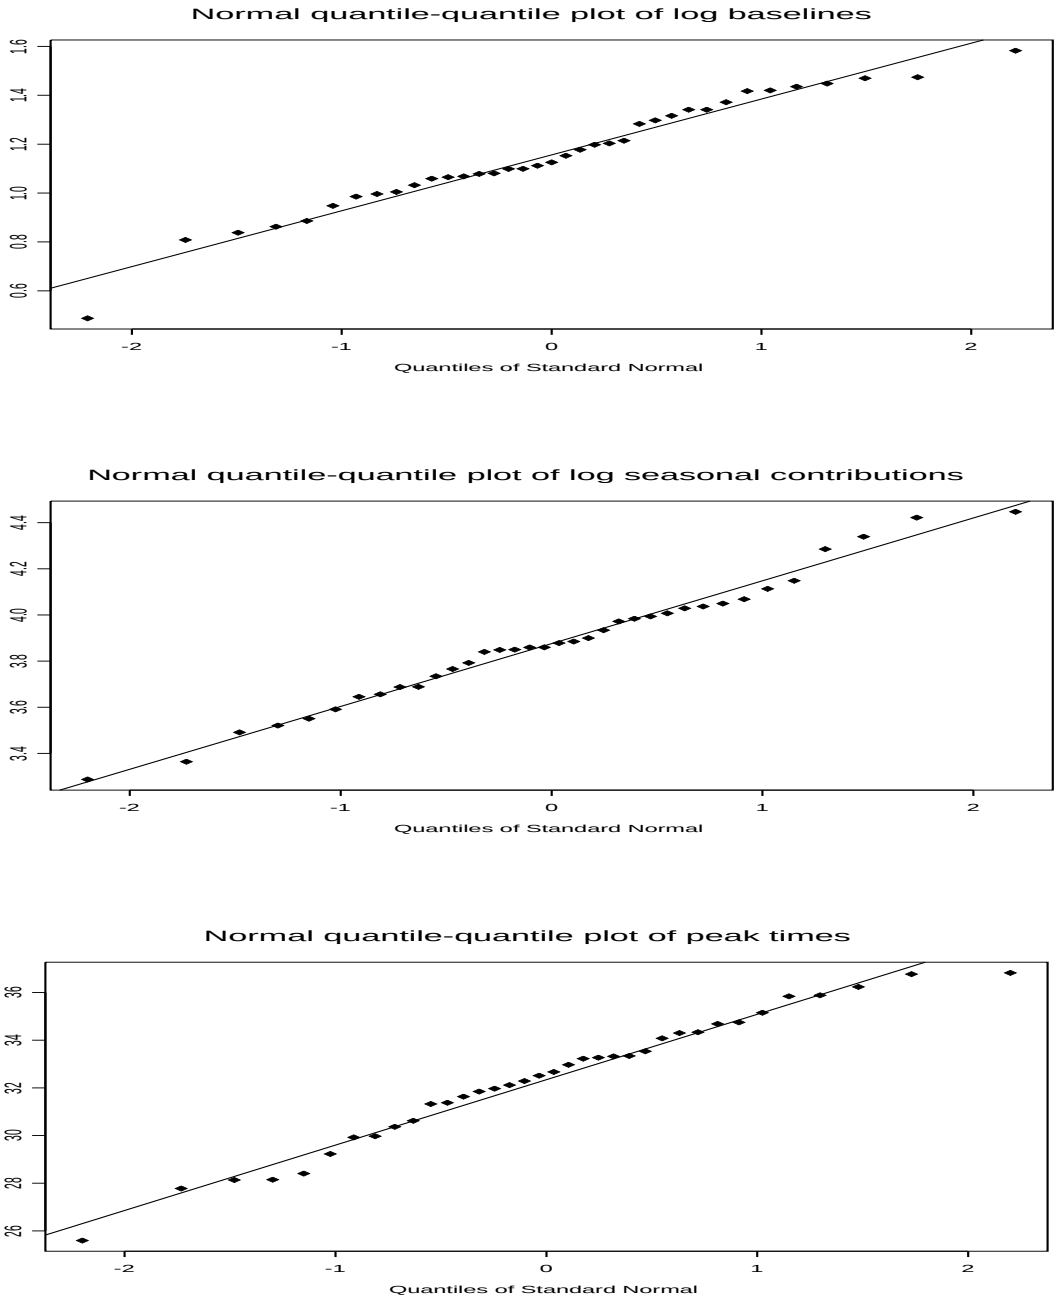

Figure 4: Normal Q-Q Plots for Estimated Baselines, Seasonal Contributions, and Times of Peak Activity

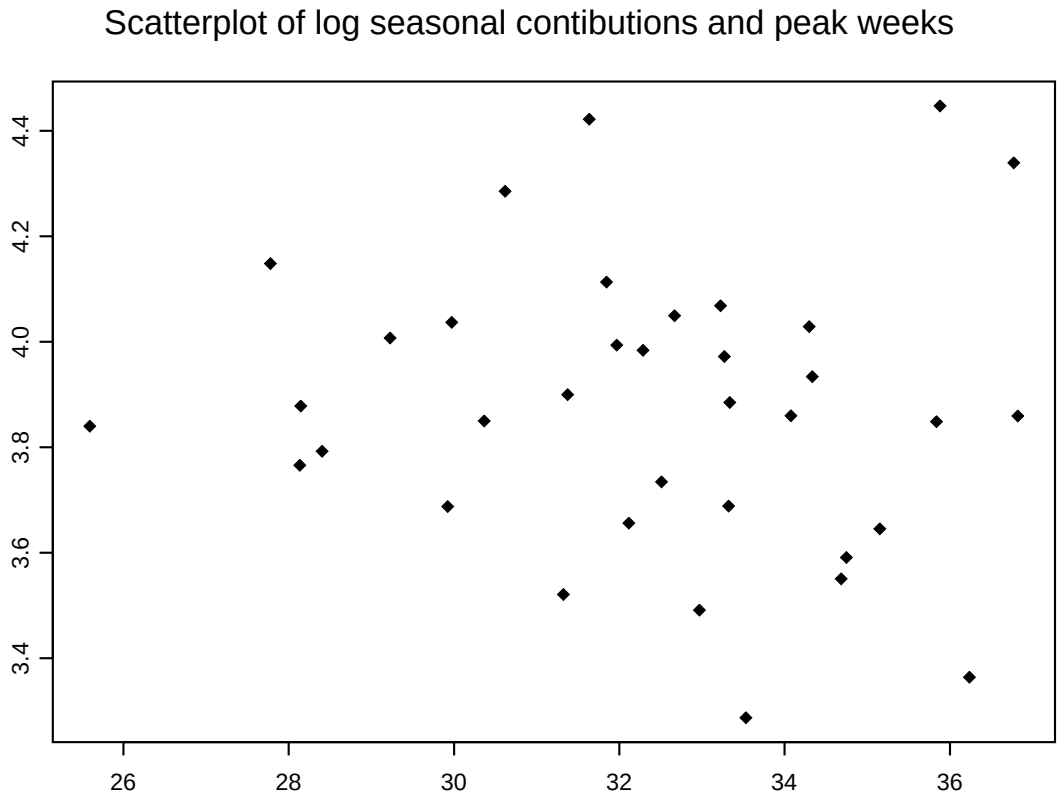

Figure 5: No Apparent Relation Between the Severity of a Flu Season and the Time of Peak Activity.

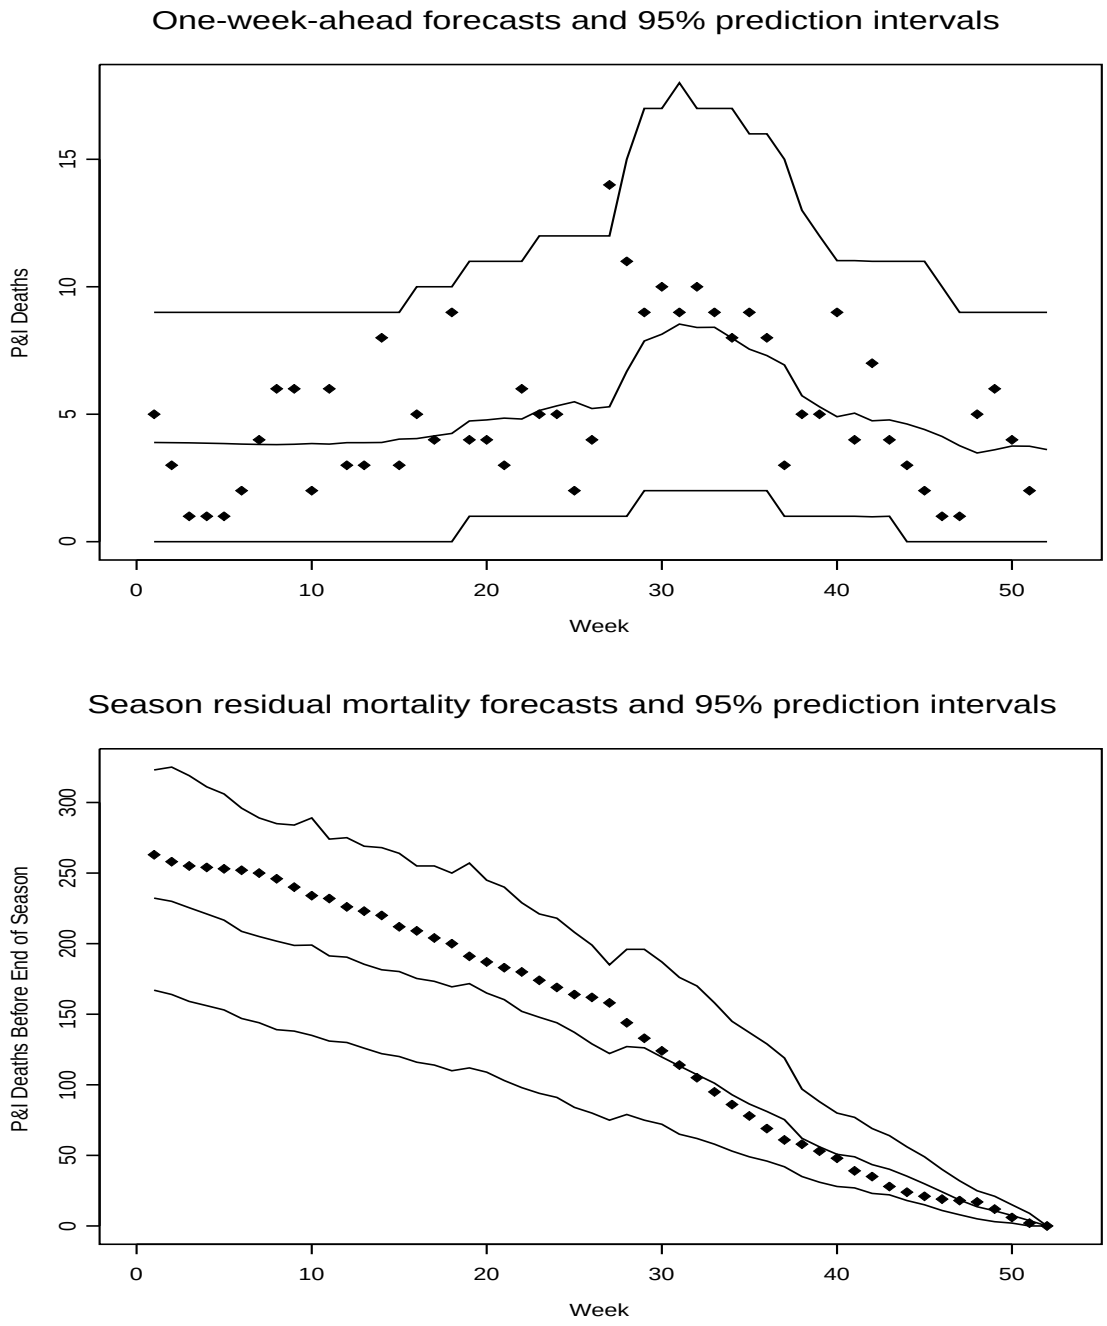

Figure 6: Predictive Ability for the Albuquerque 1998-9 Flu Season.

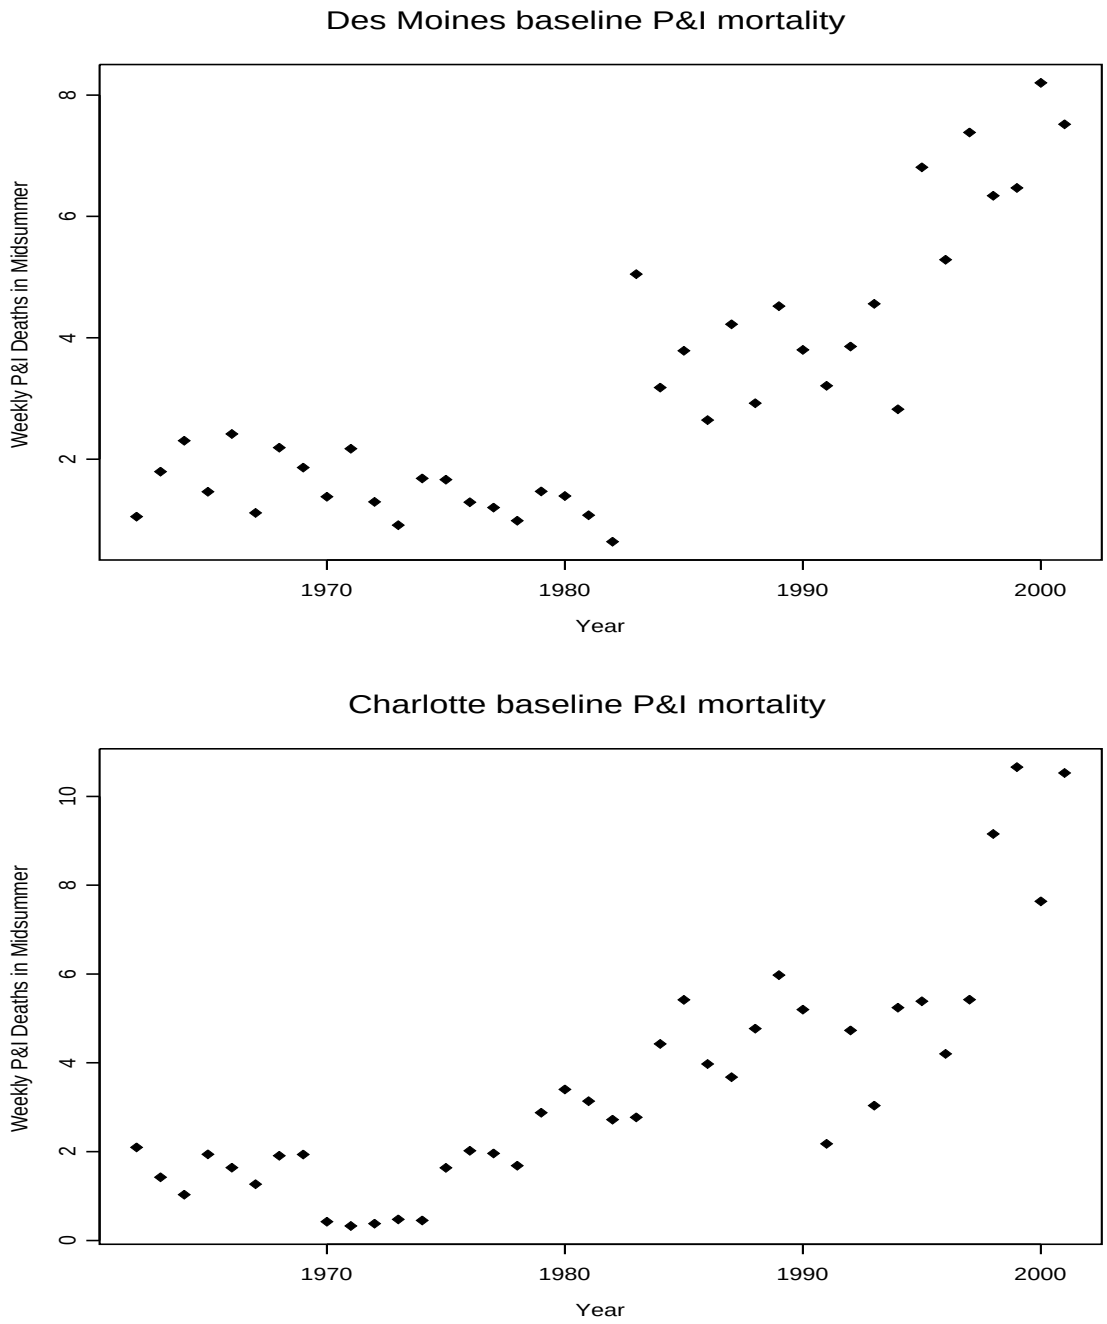

Figure 7: Baseline Behavior Not Conforming to the Steady-State Model

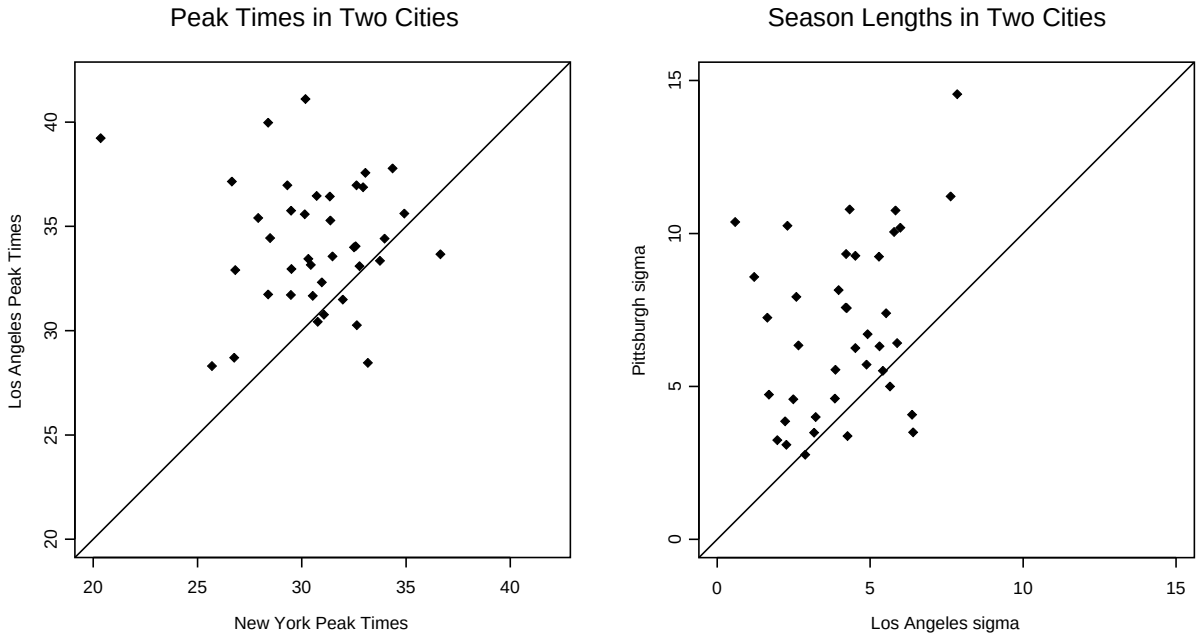

Figure 8: Estimated Times of Peak Mortality and Durations of Flu Seasons (were the cities identical for each season, all points would fall on the diagonal lines).